

PATENT COOPERATION TREATY

PCT

NOTIFICATION OF ELECTION

(PCT Rule 61.2)

From the INTERNATIONAL BUREAU

To:

Commissioner
 US Department of Commerce
 United States Patent and Trademark
 Office, PCT
 2011 South Clark Place Room
 CP2/5C24
 Arlington, VA 22202
 ETATS-UNIS D'AMERIQUE
 in its capacity as elected Office

Date of mailing (day/month/year) 28 May 2001 (28.05.01)	
International application No. PCT/GB00/03678	Applicant's or agent's file reference A25800 WO
International filing date (day/month/year) 25 September 2000 (25.09.00)	Priority date (day/month/year) 24 September 1999 (24.09.99)
Applicant HOVELL, Peter et al	

1. The designated Office is hereby notified of its election made:

☒ in the demand filed with the International Preliminary Examining Authority on:
 17 April 2001 (17.04.01)

☐ in a notice effecting later election filed with the International Bureau on:

2. The election ☒ was
☐ was not

made before the expiration of 19 months from the priority date or, where Rule 32 applies, within the time limit under Rule 32.2(b).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No.: (41-22) 740.14.35	Authorized officer Zakaria EL KHODARY Telephone No.: (41-22) 338.83.38
---	--

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
29 March 2001 (29.03.2001)

PCT

(10) International Publication Number
WO 01/22664 A1(51) International Patent Classification⁷: H04L 12/46

(21) International Application Number: PCT/GB00/03678

(22) International Filing Date:
25 September 2000 (25.09.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
99307550.6 24 September 1999 (24.09.1999) GB(71) Applicant (for all designated States except US): BRITISH
TELECOMMUNICATIONS PUBLIC LIMITED
COMPANY [GB/GB]; 81 Newgate Street, London EC1A
7AJ (GB).

(72) Inventors; and

(75) Inventors/Applicants (for US only): HOVELL, Peter

[GB/GB]; 24 Mill Road, Newbourne, Woodbridge, Suffolk, IP12 4NP (GB). KING, John, Robert [GB/GB]; 2 Hertfords Place, Chillesford, Woodbridge, Suffolk IP12 3SD (GB). PATTERSON, John [GB/GB]; The Annexe, 5 The Mills, Playford Road, Rushmere St Andrew, Ipswich, Suffolk, IP4 5RL (GB).

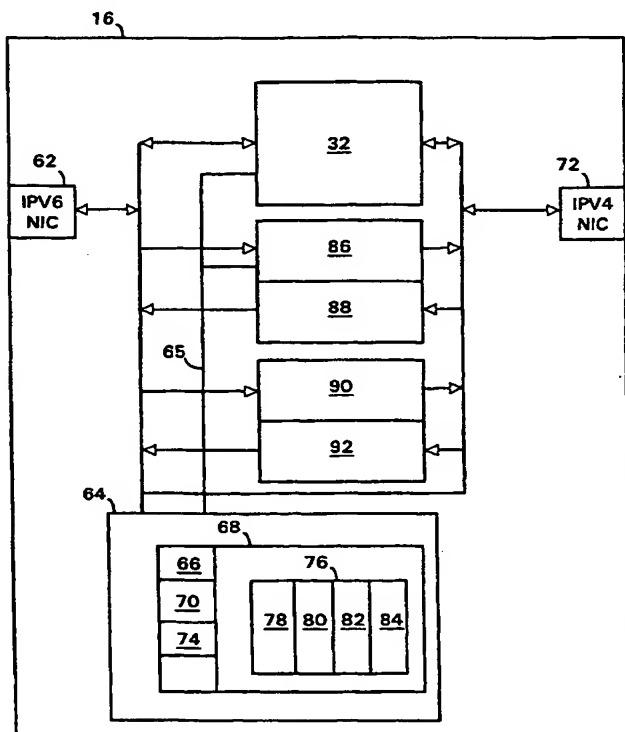
(74) Agent: SEMOS, Robert, Ernest, Vickers; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European

[Continued on next page]

(54) Title: PACKET NETWORK INTERFACING



(57) Abstract: A method of establishing a tunnel across an IPv4 domain for the transport of packets from a source host on one IPv6 domain to a destination host on another IPv6 domain, there being respective interfaces between the IPv4 domain and the IPv6 domains. The source host sends a normal IPv6 address request to its local DNS server, which relays it to an IPv6 name server in the other IPv6 domain. The response message, containing the true IPv6 address of the destination is received at the remote interface, which appends to the resulting protocol converted DNS response message a first additional record containing the true IPv6 address, and a second additional record containing the IPv4 address of that interface. Upon receipt at the near interface, the additional records are stripped off, their contents stored in an entry of a table, and the true IPv6 address written into the address record of the resulting IPv6 DNS response message. When the near interface receives a packet from an IPv6 host, it checks whether the destination address matches an entry of its table, and if so sends the packet to the encapsulator together with the IPv4 address of the remote interface. The remote interface extracts the source address and the address of the encapsulating interface and stores these in an entry in its corresponding table for use in encapsulating return packets to the source. If, however, the destination address is recognised as being of IPv4-compatible or IPv4-mapped format, the packet is sent to a protocol converter.

WO 01/22664 A1



patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— With international search report.



...PACKET NETWORK INTERFACING

This invention relates to the establishment of a tunnel entrance for packets crossing the border between a first network operating in accordance with a first transmission protocol and having network addresses in accordance with a first addressing convention, herein referred to as first type addresses, and a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention, herein referred to as second type addresses, and relates particularly, but not exclusively, to communication between hosts in respective Internet Protocol version 6 (IPv6) domains separated by an Internet Protocol version 4 (IPv4) domain.

Herein, the terms packet and message are used interchangeably and have the same meaning, and the term Internet domain is used as a specific example of a network.

In Internet technology, it had become apparent that the initial transport protocol, IPv4, needed to be enhanced, principally to increase the available address space and add a hierarchical address structure. The result was IPv6 which has a simplified header format compared with IPv4, but which uses 128 bit addresses as compared with 32 bit addresses used in IPv4.

Readers wishing to have an overview of this general area might like to access a list of Internet-Drafts, which are working documents of the IETF, at <http://www.ietf.org/1id-abstracts.txt>, and a particularly relevant document is "A Guide to the Introduction of IPv6 in the IPv4 World" <draft-ietf-ngtrans-introduction-to-ipv6-transition-01.txt>, also referred to as "Guide to IPv6 Transition".

As mentioned, the present invention relates to tunnelling. Known tunnelling techniques are of two types: configured and automatic.

A configured tunnel is created by manual configuration of a tunnelling interface between an IPv6 domain and an IPv4 domain such that all packets received from the IPv6 domain are encapsulated in IPv4 packets addressed to a specific tunnel end point, i.e. the tunnelling interface between the IPv4 domain and the remote IPv6 domain containing the destination IPv6 host.

An automatic tunnel on the other hand does not need manual configuration: the tunnel end points are automatically determined. Several automatic tunnelling

mechanisms are currently in development within the IETF, these being known in the art as, 6over4, 6to4, Dynamic Tunnelling, and Tunnel Broker.

For more detailed information on 6over4 the reader can obtain from the IETF, the document known as RFC2829, or any variant thereof, "Transmission of
5 IPv6 over IPv4 Domains without Explicit Tunnels", by B. Carpenter and C. Jung, March 1999.

For more detailed information on 6to4 the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-6to4-02.txt, or any variant thereof, "Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels", by B.
10 Carpenter and K. Moore.

For more detailed information on Dynamic Tunnelling the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-dti-00.txt.

For more detailed information on Tunnel Broker the reader can obtain from the IETF, the document known as draft-ietf-ngtrans-broker-00.txt.

15 These known automatic tunnelling mechanisms use a variety of techniques to enable the tunnel to be automatically established:

ξ 6over4 Multicast

ξ 6to4 Special IPv6 address in which the top level aggregator (TLA) contains an identifier for the 6to4 tunnelling mechanism, and the next
20 level aggregator (NLA) contains the IPv4 address of the tunnel end point

ξ Dynamic Tunnelling via DNS

ξ Tunnel Broker www based tool

In accordance with a first aspect of the present invention, there is provided a method of establishing a tunnel from a first interface between a first network and
25 a second network to a second interface between the second network and a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing
30 convention, the tunnel being for the transport of messages from a first host on the first network to a second host on the third network, the method comprising the steps of:

... sending from the first host an address request message in accordance with the first transmission protocol, referred to herein as a first type address request message, containing the name of the second host;

upon receipt of the address request message at a name to address
5 conversion system of the third network, returning an address response message in accordance with the first transmission protocol, referred to herein as a first type address response message, and containing the address of the second host in a response address field;

10 upon receipt of that first type address response message at the second interface,

converting it to an address response message in accordance with the second transmission protocol, referred to herein as a second type address response message, and

15 augmenting that converted second type address response message by fields respectively containing the address of the second interface in accordance with the second addressing convention and the address of the second host in accordance with the first addressing convention; and

20 upon receipt of that augmented that converted second type address response message at the first interface,

converting it to a first type address response message,

retrieving the contents of the augmenting fields,

25 storing at the first interface a mapping of the retrieved address of the second host and the retrieved address of the second interface for use in encapsulating messages from the first host addressed to the second host, and

replacing the content of the response address field of the resulting first type address response message by the retrieved address of the second host.

30 Preferably, for establishing a tunnel in the reverse direction for the transport of messages from the second host to the first host, the method comprises the further steps of:

upon receipt at the first interface of a message from the first host addressed to the second host, encapsulating that received message in accordance with the first mapping; and

upon receipt of the encapsulated message at the second interface,

5 un-encapsulating that received encapsulated message,
 retrieving from the encapsulating header the address of the first
 interface in accordance with the second addressing convention,
 retrieving from the un-encapsulated message the address of the first
 host in accordance with the first addressing convention, and
10 storing at the second interface a mapping of the retrieved address of
 the first host and the retrieved address of the first interface for use in
 encapsulating messages from the second host addressed to the first
 host.

 There may be included the steps of setting a time to live for a said stored
15 mapping, and rendering that stored mapping unuseable upon the expiry of the time
to live, preferably deleting that stored mapping.

 In accordance with a second aspect of the present invention, there is
provided a method of sending packets from a first host on a first network via a
second network to a second host on a third network, the first and third networks
20 operating in accordance with a first transmission protocol and having addresses in
accordance with a first addressing convention, and the second network operating in
accordance with a second transmission protocol and having addresses in accordance
with a second addressing convention, comprising the steps of:

 establishing a tunnel from a first interface between the first network and the
25 second network to a second interface between the second network and the third
network in accordance with the method of the first aspect of the present invention;

 upon receipt by the first host of the resulting first type address response
message from the first interface, retrieving the content of the response address field;

 generating one or more packets for transmission having a header including
30 source and destination address fields, the source address field containing the address
of the first host, and the destination address field containing the retrieved content of
the response address field;

 sending the or each generated packet to the first interface;

~~for the or each said~~ generated packet received by the first interface, accessing the stored mappings in accordance with the destination address of the received generated packet,

retrieving the stored interface address of the mapping whose retrieved
5 host address matches the destination address of the received generated packet,

generating an encapsulated packet having a payload formed by the received generated packet, and having a header including source and destination address fields, the source address field containing the address of the first interface, and the destination address field containing the retrieved interface network address,

10 sending the encapsulated packet to the second interface; and

for the or each encapsulated packet received by the second interface, un-encapsulating the received encapsulated packet to recover the original generated packet forming its payload, and

sending that recovered packet to the second host.

15 There may be included the step of storing the retrieved content in association with the name of the second host.

There may be included the steps of setting a time to live for the stored retrieved content, and rendering that stored retrieved content unuseable upon the expiry of the time to live, preferably deleting the stored retrieved content.

20 Protocol converters are known for enabling an IPv6 host to send messages to an IPv4 host. When a new IPv6 host is activated in an IPv6 domain, it employs the technique known as Neighbourhood Discovery (ND) to find out the identity of hosts with whom it can communicate directly. It broadcasts an ND message containing its IPv6 network address, and each host that receives it sends a reply
25 message containing that host's IPv6 network address. Since the domain uses an underlying transport mechanism, say Ethernet using media access control (MAC) addresses, each host receiving the ND message will retrieve the new host's IPv6 network address and also the MAC address of the new host, and the new host will retrieve from each reply message the sending host's IPv6 network address and its
30 MAC address.

The new host now constructs an ND table in which each entry corresponds with a neighbouring host and comprises a first part in the form of the 128 bit IPv6

address of that neighbouring host, and a second part in the form of the associated MAC address.

The interface device (containing the protocol converter) between that IPv6 domain and an adjacent IPv4 domain will also have received the ND message and sent a reply message, and the new host will have made a special Default entry in the ND table having its first part formed by 128 zeros (in variants, these are all ones) and its second part formed by the MAC address of that interface device.

Thus, when that new host wants to send a message to one of the other hosts in its domain, it constructs an IPv6 message and accesses its ND table to retrieve the MAC address associated with the destination address. The message is then encapsulated within an Ethernet packet in known manner and sent via the underlying Ethernet transport mechanism to the destination host.

If, on the other hand, the host constructs an IPv6 message having its destination address in the form of an IPv4-compatible or IPv4-mapped address, i.e. a message intended for an IPv4 host in the adjacent IPv4 domain, this destination address will not be found in the ND table. In this situation, the accessing algorithm will return the MAC address of the Default entry, and the message will be sent to the protocol converter.

Protocol converters can only convert (or translate) between corresponding fields of the headers of the IPv6 and IPv4 messages. Where a field in the header of, say, an IPv6 message has no corresponding field in the header of an IPv4 message, or vice versa, the information in that field will be lost in the protocol conversion process.

As mentioned above, tunnelling techniques are known for enabling IPv6 hosts to communicate amongst themselves when they are in spaced-apart domains. In this case, the interface device contains a tunnelling mechanism instead of a protocol converter. It will be appreciated that before now, for an IPv6 host to be able to communicate both with IPv4 hosts and with remote IPv6 hosts, it was necessary for that IPv6 host and the domain interface to be dual stack, i.e. having both IPv4 and IPv6 communication capability. If an IPv6 host was not dual stack, its accessing algorithm would return only a single MAC address for the Default entry. This would be the MAC address of the input port of a protocol converter, if the network administration had decided that the IPv6 host is to be able to communicate with

IPv4 hosts, or it would be the MAC address of the input port of a tunnelling mechanism, if the network administration had decided that the IPv6 host is to be able to communicate with IPv6 hosts. That Default entry MAC address could not be a common input address for both the protocol converter and the tunnelling
5 mechanism.

Methods of tunnel establishment in accordance with the present invention are advantageous in that the source host need only send a standard address request message in its own protocol, e.g. if the first network transmission protocol is IPv6 then that message will be an IPv6 DNS address request message, and the
10 establishment of the tunnel across the intervening IPv4 domain is automatic and completely autonomous, requiring nothing special from the source host. Thus, a company that has an existing IPv4 domain can make the transition from full IPv4 implementation to full IPv6 implementation by introducing isolated IPv6 domains using standard IPv6 technology.

15 Specific embodiments of the present invention will now be described with reference to the drawings in which:

Figure 1 is a schematic diagram of an IPv4 domain interfaced with two isolated IPv6 domains;

Figure 2 is a schematic diagram of a border router;

20 Figure 3 is a schematic diagram of an IPv6 DNS Response message;

Figure 4 is a schematic diagram of an IPv4 DNS Response message resulting from conversion of the IPv6 DNS Response message of Figure 3;

Figure 5 is a schematic diagram of an IPv6 DNS Response message resulting from conversion of the IPv4 DNS Response message of Figure 4; and

25 Figure 6 is a schematic diagram showing the format of special IPv6 addresses used with the 6to4 tunnelling technique.

In Figure 1, an IPv4 domain 10 separates a first IPv6 domain 12, constituting in accordance with the present invention a first network operating in accordance with a first transmission protocol and having network addresses in
30 accordance with a first addressing convention, from a second IPv6 domain 14, constituting in accordance with the present invention a third network also operating in accordance with the first transmission protocol and having network addresses in accordance with the first addressing convention. Hosts in the IPv4 domain 10 are



IPv4 only, and hosts in the IPv6 domains 12 and 14 are IPv6 only. The IPv4 domain 10 constitutes in accordance with the present invention a second network operating in accordance with a second transmission protocol and having network addresses in accordance with a second addressing convention. For simplifying the drawings, no 5 IPv4 hosts are shown, and in each of the IPv6 domains 12 and 14 only one IPv6 host (28 and 30, respectively, as referred to later) is shown.

The first IPv6 domain 12 is connected to the IPv4 domain 10 via a border router 16A, also referred to as an ingress interface with respect to the first IPv6 domain 12 and constituting an interface of the present invention. The second IPv6 10 domain 14 is connected to the IPv4 domain 10 via a border router 16B, also referred to as an egress interface with respect to the first IPv6 domain 12. The border router 16B is identical to the border router 16A.

The IPv4 domain 10 contains a complete domain name system (DNS) 20 including a plurality of DNS servers 22, of which only two DNS servers 22A and 22B 15 are shown, and the IPv6 domains 12 and 14 contain respective DNS servers 24 and 26.

Suppose that a host 28 in the first IPv6 domain 12 wishes to send a packet to a host 30 in the second IPv6 domain 14. Thus, for this transaction, the host 28 is referred to as the source host 28 and the host 30 is referred to as the destination 20 host 30.

The source host 28 knows the name of the destination host 30, so it constructs in known manner an IPv6 DNS Request message (not shown) requesting the IPv6 address of the destination host 30. The source host 28 sends this DNS Request message as a recursive request to its local DNS server, which in this 25 embodiment is the DNS server 24. The DNS server 24 will, in known manner, send a number of iterative DNS Request messages (not shown) to the DNS 20 until it learns about the DNS server 26. Finally, a DNS Request message (not shown) will go to the DNS server 26 requesting the IPv6 address of the destination host 30. As the DNS request passes from the first IPv6 domain 12 through the border router 16A to the 30 IPv4 domain 10, it is processed by a protocol converter (PC) 32A (see Figure 2) and undergoes IPv6/IPv4 translation. Correspondingly, as the DNS request passes from the IPv4 domain 10 through the border router 16B to the second IPv6 domain 14, it is processed by a protocol converter 32B and undergoes IPv4/IPv6 translation.



The protocol converters 32A and 32B conform to a specification known as Network Address Translation-Protocol Translation (NAT-PT). They translate between IPv4 and IPv6 addresses and keep state during the time of the session, and translation between IPv4 and IPv6 DNS Request messages and DNS Response messages, including translation of IP headers and DNS payload information is controlled by an application layer gateway (ALG). In the art, an alternative term for a DNS Response message is a DNS Answer message.

For more detailed information the reader can obtain from the Internet Engineering Task Force (IETF), the document known as draft-ietf-ngtrans-natpt-05.txt, or any variant thereof, "Network Address Translation - Protocol Translation (NATPT)", by G. Tsirtsis and P. Srishuresh.

The DNS server 26 responds to the DNS Request message for the IPv6 address of the destination host 30 with a IPv6 DNS Response message 34 (see Figure 3) having a conventional format of destination address field 36, source address field 38, and response address record 40 that contains the IPv6 address of the destination host 30.

This IPv6 DNS Response message 34 travels through the second IPv6 domain 14 to the border router 16B where it becomes converted to an IPv4 DNS Response message 42 (see Figure 4) comprising destination address field 44, source address field 46, response address record 48 and, in accordance with the present invention, additional records 50 and 52, and this message 42 travels through the IPv4 domain 10 to the border router 16A where it becomes converted to an IPv6 DNS Response message 54 comprising destination address field 56, source address field 58 and response address record 60, and sent to the source host 28. The route that the DNS Response message (in its various forms, i.e. 34, 42 and 54) takes in the domains 10, 12 and 14 depends on the DNS configuration in each of the domains, but it will have to pass through the border router 16B and the border router 16A in that order.

For convenience, the terms "field" and "record" are used synonymously and interchangeably in this description, although in the art a field is generally taken to be a component part of a record.

When the IPv6 DNS Response message 34 is received by the border router 16B via its IPv6 network interface controller 62B, it is fed in parallel to the protocol converter 32B and to a controller 64B, and also to an encapsulator 86B and a 6to4

encapsulator 90B. The controller 64B is connected via a control line 65A to control inputs of the protocol converter 32B, the encapsulator 86B and the 6to4 encapsulator 90B and by placing a suitable address on the control line 65A selects the appropriate one of these devices.

5 The controller 64B (i) recognises that the received message is a DNS Response message and enables the protocol converter 32B, (ii) retrieves the IPv6 address from the response record 40 and writes this message to a storage location 66B formed by part of its internal memory 68B; (iii) writes the IPv4 address of the border router 16B, i.e. the IPv4 address of the tunnel terminating endpoint on the
10 border router 16B, to a storage location 70B also formed by part of its internal memory 68B; (iv) receives from the protocol converter 32B the converted DNS Response message 42; appends the content of the storage location 70B as the first additional record 50, and the content of the storage location 68B as the second additional record 52; and (v) sends the resulting IPv4 DNS Response message 42 to
15 an IPv4 network interface controller 72B for transmission over the IPv4 domain 10 to the border router 16A.

 In a variant, the received IPv6 DNS Response message 34 is fed only to the controller 64B, which writes this message to a storage location 74B formed by part of its internal memory 68B. The controller 64B then (i) retrieves the IPv6 address
20 from the response record 40 and writes this message to the storage location 66B; (ii) writes the IPv4 address of the border router 16B to the storage location 70B; (iii) generates a modified IPv6 DNS Response message by retrieving the content of the storage location 66B and appending the content of the storage location 70B as the first additional record 50, and the content of the storage location 66B as the second
25 additional record 52; and (iv) sends this modified IPv6 DNS Response message to the protocol converter 32B. The ALG of the protocol converter 32 processes only the header and address response record of the DNS Response message to produce the resulting IPv4 DNS Response message 42, i.e. it allows the additional records to remain unchanged.

30 It will be appreciated that the feeding of the received message directly to the protocol converter 32B, and the sending of an enabling signal on the control line 65A to the protocol converter 32B by the controller 64B is logically equivalent to the sending of the received message to the protocol converter 32B by the controller 64B



in the above variant, and constitutes sending the message to the protocol converter 32B in accordance with the present invention.

When the IPv4 DNS Response message 42 is received by the border router 16A via its IPv4 network interface controller 72A, it is fed in parallel to the protocol converter 32A and to a controller 64A.

The controller 64A (i) receives from the protocol converter 32A the output IPv6 DNS Response message comprising destination address field 56, source address field 58, response address record 60, and additional records 50, 52; and (ii) retrieves the IPv6 address from the second additional record 52, i.e. the true IPv6 address of the destination host 30, and inserts it into the response address record 60 of that output message instead of the IPv4-compatible IPv6 address for the destination host 30, which the protocol converter 32A had generated. The controller 64A then strips off the additional records 50, 52, and sends the resulting IPv6 DNS Response message 54 (see Figure 5) to an IPv6 network interface controller 62A of the border router 16A for transmission to the source host 28.

Additionally, the controller 64A is arranged to retrieve the IPv4 address of the tunnel terminating endpoint from the first additional record 50, to create a mapping of the IPv6 address of the destination host 30 to the address of the IPv4 tunnel terminating endpoint, and to store this mapping, i.e. create an entry, in an IPv6/tunnel endpoint table 76A formed by part of an internal memory 68A of the controller 64A. This table is referred to as a look-up table.

In a variant, the additional records 50, 52 are stripped from the DNS Response message prior to the retrieval of their contents. In another variant, the additional records remain attached to the DNS Response message, but this is not as efficient as stripping the additional records. In this present embodiment, each entry in the IPv6/tunnel endpoint table 76A comprises a first element 78A comprising the IPv6 address of a corresponding destination host 30, a second element 80A comprising the IPv4 address of the tunnel terminating endpoint, i.e. of the border router 16B, and third and fourth elements 82A and 84A, to be described later.

Upon receipt of the resultant IPv6 DNS Response message 54, the source host 28 retrieves the IPv6 address from its address record 60 and stores it in its internal memory for use in sending data packets to the destination host 30.

In known manner, the source host 28 generates for each of these data packets a header including a source address field and a destination address field, and writes the retrieved IPv6 address into the destination address field.

Upon receipt of each of these data packets at the border router 16A, the controller 64A retrieves the destination address, and, in accordance with the retrieved destination address, accesses the IPv6/tunnel endpoint table 76A. If there is a match with the contents of a first element 78A of an entry, the controller 64A retrieves the corresponding IPv4 tunnel terminating endpoint from the second element 80A of that entry, and commands an encapsulator 86A to encapsulate the packet in an IPv4 packet addressed to the border router 16B using the retrieved IPv4 tunnel terminating endpoint address. In this embodiment, the encapsulator 86A is arranged to receive the packet directly from the IPv6 network interface controller 62A of the border router 16A, but it does not perform encapsulation until commanded by the controller 64A. In a variant, the controller 64A receives the packet directly from the IPv6 network interface controller 62A and passes it to the encapsulator 86A if there is a match. In practice, when the border router 16A receives a packet the controller 64A writes it into a storage location of its internal memory 68A, and the controller 64A will pass to the encapsulator 86A the address of the relevant storage location, together with an instruction for the encapsulator 86A to access that storage location.

Upon receipt of the encapsulating IPv4 packet at the border router 16B, an un-encapsulator 88B of the border router 16B strips off the IPv4 header and retrieves the payload of that IPv4 packet, i.e. un-encapsulates the original IPv6 packet from the source host 28, and sends that IPv6 packet to the destination host 30. The controller 64B also creates a mapping (in its IPv6/tunnel endpoint table 76B) of the IPv6 address of the source host 28 and the tunnel originating endpoint, i.e. the IPv4 address of the originating border router 16A, these being respectively retrieved from the source address field of the IPv6 header and the source address field of the IPv4 header.

When the destination host 30 returns a Reply packet, a controller 64B of the border router 16B retrieves the destination address, "IPv6 host 28", from the received Reply packet, accesses its IPv6/tunnel endpoint table 76B in accordance with the retrieved destination address (i.e. seeking a matching element 78B),



retrieves the corresponding IPv4 address (element 80B), and commands an encapsulator 86B to encapsulate the Reply packet in an IPv4 packet addressed to an un-encapsulator 88A of border router 16A using the IPv4 tunnel originating endpoint address just retrieved from the element 80B of the IPv6/tunnel endpoint table 76B.

- 5 Upon receipt of this IPv4 packet at the border router 16A, the un-encapsulator 88A performs un-encapsulation to retrieve the Reply packet, and the border router 16A then sends the retrieved Reply packet to the source host 28.

The source host 28 and the destination host 30 are now in a session in which IPv6 packets pass between them via the tunnel just established between the
10 border routers 16A and 16B.

The above described mechanism provides for an IPv6 host, which is in an isolated IPv6 domain, to communicate with another IPv6 host, which is in another isolated IPv6 domain, via an intermediate IPv4 domain, without any knowledge of where that other IPv6 host is, and without the source IPv6 host needing to do
15 anything different from a standard communication procedure with another IPv6 host within its own IPv6 domain. The DNS server local to the source IPv6 host makes a Request via the IPv4 domain to the IPv6 DNS server that is on the same network as the destination IPv6 host, and the border routers automatically set up respective mappings associating the tunnel endpoint and the IPv6 address of the IPv6 hosts
20 behind the border routers.

In this embodiment, the host 28 stores the address retrieved from the DNS Response for a short time. The host 28 records the time at which the address is stored and deletes that stored address twenty four hours later. Alternatively, other methods can be employed, e.g. the address can be deleted at midnight, when the
25 day changes, or the address can be stored with a "time to live" value which is the start value of a count down timer. Similarly, the border router 16A deletes the entry in its IPv6/tunnel endpoint table 76A after a short time. By preventing permanent storage of the address of the host 30, this avoids the risk that the tunnel terminating endpoint becomes out of date for transmissions to the host 30, e.g. the border
30 router 16B is out of service, and that the border router 16A continues to send encapsulated packets to that out of date address. In a variant, the address is deleted after transmission of the last packet of the session. The intent of all the above



alternatives is, to a greater or lesser extent, to ensure that tunnels across the IPV4 domain 10 are always up to date.

In alternative embodiments, some of the entries in the IPv6/tunnel endpoint table 76A can be created by the administrative personnel of the network operator.

- 5 This is known as manual configuration of a tunnel, and the tunnel is permanent until changed at a later date by the administrative personnel.

As shown in Figure 2, the border router 16A also comprises a 6to4 tunnelling encapsulator 90A (and 6to4 tunnelling un-encapsulator 92A) and can thus interwork with a border router which is similarly enabled, although in variants these
10 may be omitted. The special IPv6 addresses 94 (see Figure 6) used for this technique have a three-part format of which a first part 96 having thirty two bits is a prefix uniquely identifying that the packet is to be tunnelled by the 6to4 tunnelling technique, a second part 98 having thirty two bits is the IPv4 address of the 6to4 tunnel endpoint, and the third part 100 having sixty four bits is known as the
15 interface ID which is the modified MAC address of the destination host. In variants having a different tunnelling encapsulator, a different respective prefix is used for the same purpose.

In some variants of these embodiments, the controller 64A is arranged to recognise the presence of this prefix within the retrieved destination address of a
20 received packet and to command the 6to4 tunnelling encapsulator 90A to handle the received packet, and in this case the 6to4 tunnelling encapsulator 90A is arranged to retrieve the special IPv6 address and to extract from its second part 98 the IPv4 address of the 6to4 tunnel endpoint.

Where, as mentioned above, the controller 64A is arranged to perform prefix
25 recognition, the prefix to be recognised is stored in a storage location of its internal memory 68A, and this storage location can be an entry or part of an entry of the IPv6/tunnel endpoint table 76A.

The un-encapsulators 88B and 92B have respective IPv4 addresses, which are used by the corresponding encapsulators 86A and 90A in generating their
30 respective encapsulated packets.

In the preferred arrangement of border router having a plurality of different encapsulators, e.g. 86, 90, the controller 64A accesses the IPv6/tunnel endpoint

table 76A in accordance with a set of match criteria to cover the range of possible situations. These are

(a) a tunnel has already been established by the above described DNS Request technique and a IPv6 destination-specific IPv6/IPv4 entry exists in the
5 IPv6/tunnel endpoint table 76A;

(b) a tunnel has already been established by one of the known tunnelling techniques and an IPv6/IPv4 entry exists in the IPv6/tunnel endpoint table 76A, of which the first element 78A has a first part in the form of a specific prefix corresponding to that tunnelling technique;

10 (c) network management personnel have manually configured the border router 16 to define a tunnel to a specific IPv6 destination host using 6to4 (or 6over4) to another border router (which might be the border router 16B or a different border router (not shown) associated with a further IPv6 domain (not shown)), and for this case the IPv6/tunnel endpoint table 76A has an entry whose first element
15 78A is the IPv4-compatible (or IPv4-mapped) address of that destination host;

(d) network management personnel have manually configured the border router 16A to define a tunnel to unspecified IPv6 destination hosts using 6to4 (or 6over4) to another border router (which might be the border router 16B or a different border router associated with a further IPv6 domain), and for this case the
20 IPv6/tunnel endpoint table 76A has an entry having its first element 78A in the form of the 6to4 (or 6over4) prefix followed by the IPv4 address of that other border router followed by null characters, and in some variants the second element 80A of this entry contains null characters, whilst in yet other variants the second element 80A of this entry contains the IPv4 address of that other border router; and

25 (e) the table contains an entry whose first element 78A is a generic IPv4-compatible or IPv4-mapped IPv6 address, i.e. its first eighty bits are all zeros and the final thirty two (or in a variant, forty eighty) bits are null characters (zeros), the second and fourth elements 80A and 84A contain null characters, and a third element 82A contains the identifier "PC".

30 The controller 64A uses its IPv6/tunnel endpoint table 76A to determine the appropriate handling of a received packet in the following manner.

If the controller 64A finds an entry having a first element 78A matching the complete retrieved destination address, then the content of the second element 80A

of that entry is retrieved and used as the IPv4 destination address, i.e. that of the border router 16B, the tunnel endpoint. Additionally, the content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the encapsulator 86A. The content of the third element 82A is either an identifier for an encapsulator 86A, 90A (e.g. "EN") or an identifier for the protocol converter 32A (e.g. "PC"). As a further check, the fourth element 84A of that entry contains an identifier for the encapsulation type. In other words, for an entry matching the complete retrieved destination address, as just described, the encapsulation type identifier will be "DNS" to signify that the encapsulator 86A is to be used.

If the controller 64A finds an entry whose first element 78A has the first thirty two bits matching the first thirty two bits, i.e. the special 6to4 prefix part, of the retrieved destination address, then the third and fourth elements 82A and 84A of that entry are checked ("EN" and "6to4", respectively), the IPv4 destination address is retrieved from the second element 80A of that entry and sent with the packet received by the border router 16A to be processed by the encapsulator 90A.

If the retrieved destination address is either IPv4-compatible or IPv4-mapped, i.e. the packet is to be protocol converted for an IPv4 destination, then its first eighty bits will be all zeros, and the following sixteen bits will be either all zeros if the address is IPv4-compatible, or all ones if the address is IPv4-mapped. Thus the controller 64A checks to see whether its IPv6/tunnel endpoint table 76A contains an entry whose first element 78A has its first eighty bits all zeros. The content of the second element 80A of such an entry will be all null characters, because no tunnel is involved, and the content of the fourth element 84A of such an entry will be all null characters, because no encapsulation is involved. The content of the third element 82A of that entry is retrieved and used as a check that the retrieved IPv4 destination address and the packet received by the border router 16A are to be processed by the protocol converter 32A. In this case, the content of the third element 82A is an identifier for the protocol converter 32A (e.g. "PC").

In other variants, the controller 64A is arranged to access the IPv6/tunnel endpoint table 76A, as before, and only command the 6to4 tunnelling encapsulator 90A upon detecting a match with an entry, and in this case the controller 64A passes the special IPv6 address to the 6to4 tunnelling encapsulator 90A, or

alternatively the controller 64A extracts the IPv4 address of the 6to4 tunnel endpoint and passes it to the 6to4 tunnelling encapsulator 82A, or yet again the controller 64A, upon detecting such a match, retrieves the element 80A of the matching entry. This element 80A contains the IPv4 address of the 6to4 tunnel endpoint, which was inserted into that element 80A by the controller (or manually) upon creation of that entry.

In the above embodiment, the local DNS server for the source host 28 is the IPv6 DNS server 24, but in alternative embodiments it can be one of the IPv4 servers 22 of the DNS 20. In such alternative embodiments, although the host 28 can send a DNS Request message for obtaining the IPv6 address of the host 30 and, by the present invention, establish a tunnel across the IPv4 domain, the situation is not symmetrical in that the host 30 cannot act as a source and establish a corresponding tunnel across the IPv4 domain to the host 28. For an IPv6 host to be contactable, i.e. act as destination, by the method of the present invention, it is necessary for the DNS server local to that IPv6 host to be an IPv6 DNS server on the same IPv6 domain as that IPv6 host, because the DNS Response message has to pass through the border router adjacent to that IPv6 host in order that the tunnel can be established. In other words, the DNS Request message has to pass through the IPv4 domain and not stop at an IPv4 DNS server acting as the local DNS server for the intended destination IPv6 host.

CLAIMS

1. A method of establishing a tunnel from a first interface between a first network and a second network to a second interface between the second network and a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, the tunnel being for the transport of messages from a first host on the first network to a second host on the third network, the method comprising the steps of:

sending from the first host an address request message in accordance with the first transmission protocol, referred to herein as a first type address request message, containing the name of the second host;

upon receipt of the address request message at a name to address conversion system of the third network, returning an address response message in accordance with the first transmission protocol, referred to herein as a first type address response message, and containing the address of the second host in a response address field;

upon receipt of that first type address response message at the second interface,

converting it to an address response message in accordance with the second transmission protocol, referred to herein as a second type address response message, and

augmenting that converted second type address response message by fields respectively containing the address of the second interface in accordance with the second addressing convention and the address of the second host in accordance with the first addressing convention; and

upon receipt of that augmented that converted second type address response message at the first interface,

converting it to a first type address response message,
retrieving the contents of the augmenting fields,

- storing at the first interface a mapping of the retrieved address of the second host and the retrieved address of the second interface for use in encapsulating messages from the first host addressed to the second host, and

5 replacing the content of the response address field of the resulting first type address response message by the retrieved address of the second host.

2. A method as claimed in claim 1, wherein for establishing a tunnel in the reverse direction for the transport of messages from the second host to the first host, the method comprises the further steps of:

upon receipt at the first interface of a message from the first host addressed to the second host, encapsulating that received message in accordance with the first mapping; and

15 upon receipt of the encapsulated message at the second interface, un-encapsulating that received encapsulated message, retrieving from the encapsulating header the address of the first interface in accordance with the second addressing convention, retrieving from the un-encapsulated message the address of the first host in accordance with the first addressing convention, and

20 storing at the second interface a mapping of the retrieved address of the first host and the retrieved address of the first interface for use in encapsulating messages from the second host addressed to the first host.

25

3. A method as claimed in either claim 1 or claim 2, including setting a time to live for a said stored mapping, and rendering that stored mapping unuseable upon the expiry of the time to live.

30 4. A method as claimed in claim 3, wherein the rendering step deletes that stored mapping.



5. A method of sending packets from a first host on a first network via a second network to a second host on a third network, the first and third networks operating in accordance with a first transmission protocol and having addresses in accordance with a first addressing convention, and the second network operating in accordance with a second transmission protocol and having addresses in accordance with a second addressing convention, comprising the steps of:

establishing a tunnel from a first interface between the first network and the second network to a second interface between the second network and the third network in accordance with the method of claim 1;

upon receipt by the first host of the resulting first type address response message from the first interface, retrieving the content of the response address field;

generating one or more packets for transmission having a header including source and destination address fields, the source address field containing the address of the first host, and the destination address field containing the retrieved content of

the response address field;

sending the or each generated packet to the first interface;

for the or each said generated packet received by the first interface, accessing the stored mappings in accordance with the destination address of the received generated packet,

retrieving the stored interface address of the mapping whose retrieved host address matches the destination address of the received generated packet,

generating an encapsulated packet having a payload formed by the received generated packet, and having a header including source and destination address fields, the source address field containing the address of the first interface,

and the destination address field containing the retrieved interface network address,

sending the encapsulated packet to the second interface; and

for the or each encapsulated packet received by the second interface, un-encapsulating the received encapsulated packet to recover the original generated packet forming its payload, and

sending that recovered packet to the second host.

6. A method as claimed in claim 5, including the step of storing the retrieved content in association with the name of the second host.

7. A method as claimed in claim 6, including the steps of setting a time to live for the stored retrieved content, and rendering that stored retrieved content unuseable upon the expiry of the time to live.

5

8. A method as claimed in claim 7, wherein the rendering step deletes the stored retrieved content.

9. A method of establishing a tunnel from a first interface between a first
10 network and a second network to a second interface between the second network and a third network, the method being substantially as hereinbefore described with reference to the drawings.

10. A method of sending packets from a first host on a first network via a
15 second network to a second host on a third network, the method being substantially as hereinbefore described with reference to the drawings.

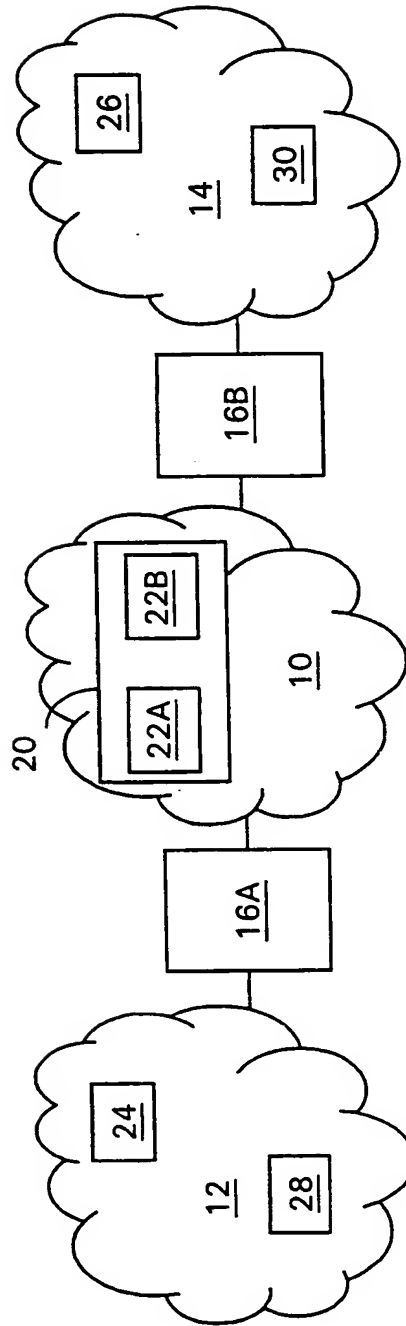


Fig.1

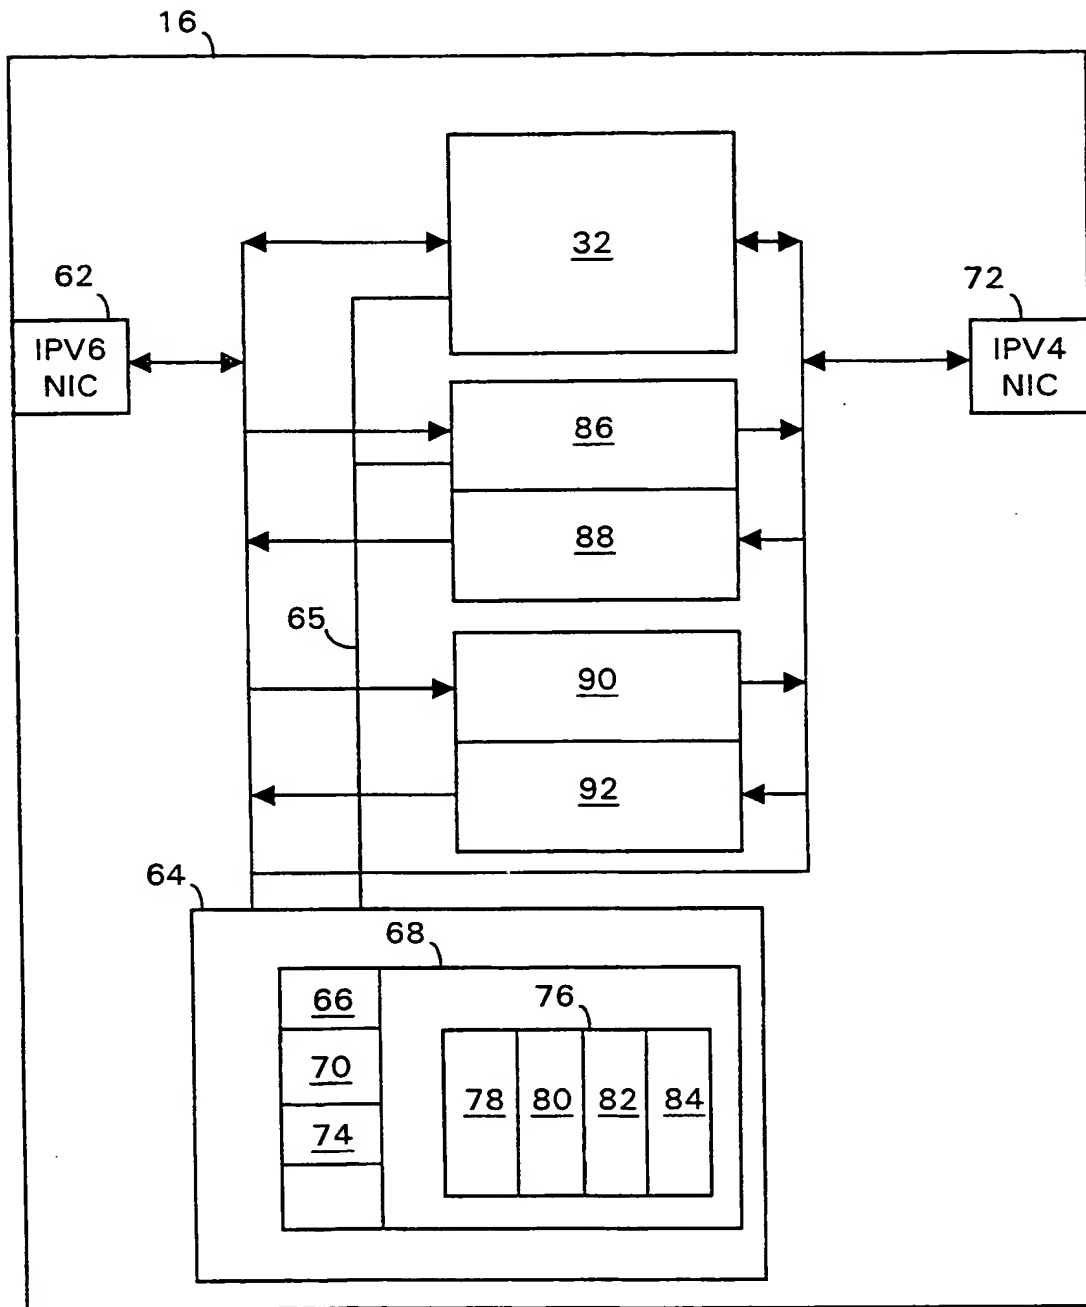


Fig. 2

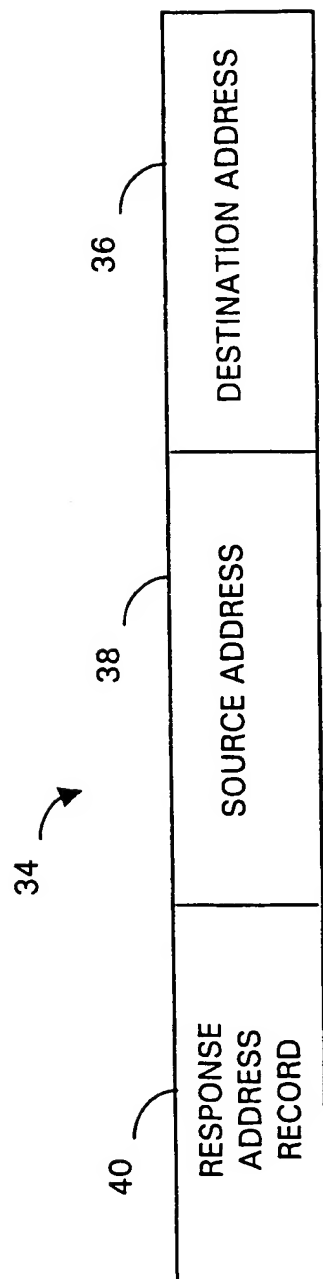


Fig. 3

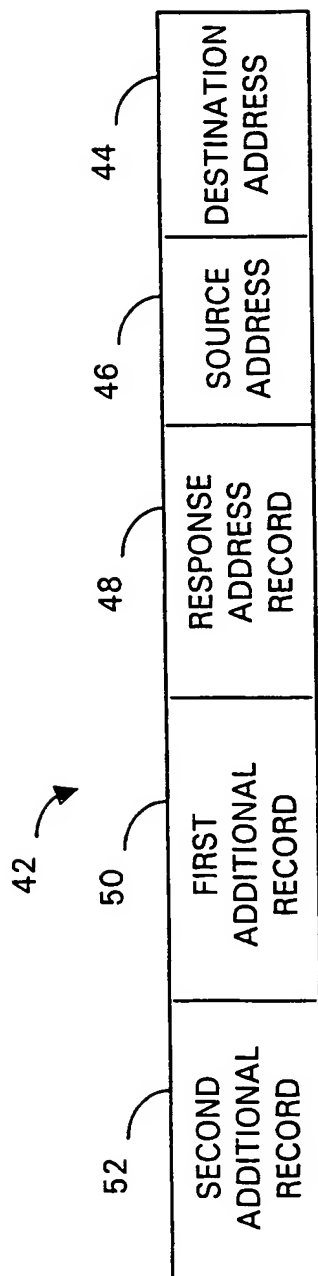


Fig. 4

5/6

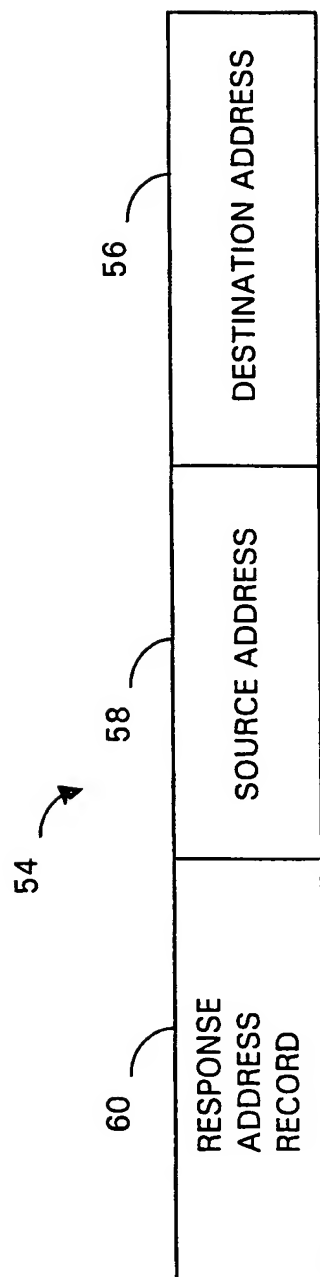


Fig. 5

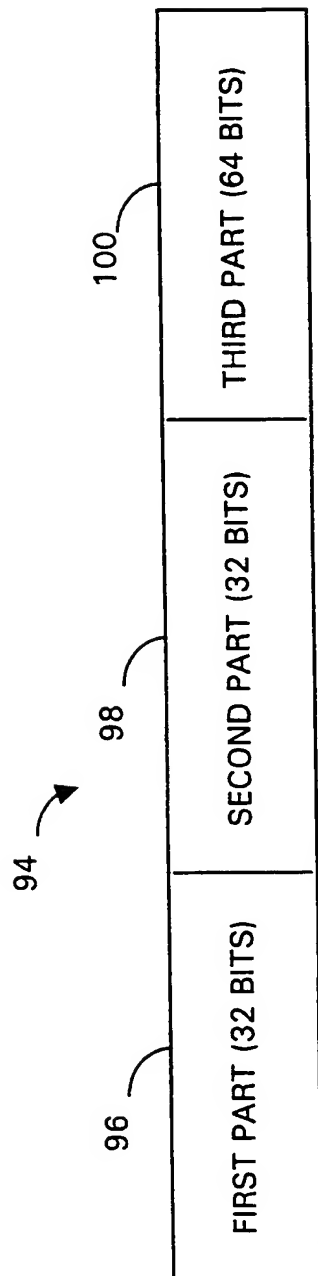


Fig. 6

INTERNATIONAL SEARCH REPORT

International Application No
PCT/GB 00/03678

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L12/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

PAJ, EPO-Internal, IBM-TDB, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 802 285 A (HIRVINIEMI SEPPÖ) 1 September 1998 (1998-09-01) abstract page 2, right-hand column, line 13 -page 3, left-hand column, line 28	1-10
A	LABORDE D: "UNDERSTANDING AND IMPLEMENTING EFFECTIVE VPNS" COMPUTER TECHNOLOGY REVIEW,US,WESTWORLD PRODUCTION CO. LOS ANGELES, vol. 18, no. 2, 1 February 1998 (1998-02-01), page 12,14,16 XP000733729 ISSN: 0278-9647 the whole document	1-10



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *S* document member of the same patent family

Date of the actual completion of the international search

20 November 2000

Date of mailing of the international search report

28/11/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 00/03678

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 840 482 A (HITACHI LTD) 6 May 1998 (1998-05-06) claims -----	1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 00/03678

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5802285 A	01-09-1998	FI 90710 B GB 2267418 A, B	30-11-1993 01-12-1993
EP 0840482 A	06-05-1998	JP 10136052 A JP 11055319 A US 6118784 A	22-05-1998 26-02-1999 12-09-2000

PCT

REC'D 24 JAN 2002

INTERNATIONAL PRELIMINARY EXAMINATION REPORT PCT

(PCT Article 36 and Rule 70)

Applicant's or agent's file reference A25800 WO	FOR FURTHER ACTION See Notification of Transmittal of International Preliminary Examination Report (Form PCT/IPEA/416)	
International application No. PCT/GB00/03678	International filing date (day/month/year) 25/09/2000	Priority date (day/month/year) 24/09/1999
International Patent Classification (IPC) or national classification and IPC H04L12/46		
Applicant BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY		

1. This international preliminary examination report has been prepared by this International Preliminary Examining Authority and is transmitted to the applicant according to Article 36.


2. This REPORT consists of a total of 11 sheets, including this cover sheet.

- ☐ This report is also accompanied by ANNEXES, i.e. sheets of the description, claims and/or drawings which have been amended and are the basis for this report and/or sheets containing rectifications made before this Authority (see Rule 70.16 and Section 607 of the Administrative Instructions under the PCT).

These annexes consist of a total of sheets.

3. This report contains indications relating to the following items:

- I ☒ Basis of the report
- II ☐ Priority
- III ☒ Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- IV ☐ Lack of unity of invention
- V ☒ Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- VI ☐ Certain documents cited
- VII ☐ Certain defects in the international application
- VIII ☐ Certain observations on the international application

Date of submission of the demand 17/04/2001	Date of completion of this report 21.01.2002
Name and mailing address of the international preliminary examining authority:  European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Tx: 523656 epmu d Fax: +49 89 2399 - 4465	Authorized officer Keller, M Telephone No. +49 89 2399 8807



**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB00/03678

I. Basis of the report

1. With regard to the **elements** of the international application (*Replacement sheets which have been furnished to the receiving Office in response to an invitation under Article 14 are referred to in this report as "originally filed" and are not annexed to this report since they do not contain amendments (Rules 70.16 and 70.17)*):

Description, pages:

1-17 as originally filed

Claims, No.:

1-10 as originally filed

Drawings, sheets:

1/6-6/6 as originally filed

2. With regard to the **language**, all the elements marked above were available or furnished to this Authority in the language in which the international application was filed, unless otherwise indicated under this item.

These elements were available or furnished to this Authority in the following language: , which is:

- ☐ the language of a translation furnished for the purposes of the international search (under Rule 23.1(b)).
- ☐ the language of publication of the international application (under Rule 48.3(b)).
- ☐ the language of a translation furnished for the purposes of international preliminary examination (under Rule 55.2 and/or 55.3).

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, the international preliminary examination was carried out on the basis of the sequence listing:

- ☐ contained in the international application in written form.
- ☐ filed together with the international application in computer readable form.
- ☐ furnished subsequently to this Authority in written form.
- ☐ furnished subsequently to this Authority in computer readable form.
- ☐ The statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.
- ☐ The statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished.

4. The amendments have resulted in the cancellation of:

- ☐ the description, pages:
- ☐ the claims, Nos.:

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB00/03678

☐ the drawings, sheets:

5. ☐ This report has been established as if (some of) the amendments had not been made, since they have been considered to go beyond the disclosure as filed (Rule 70.2(c)):

(Any replacement sheet containing such amendments must be referred to under item 1 and annexed to this report.)

6. Additional observations, if necessary:

III. Non-establishment of opinion with regard to novelty, inventive step and industrial applicability

1. The questions whether the claimed invention appears to be novel, to involve an inventive step (to be non-obvious), or to be industrially applicable have not been examined in respect of:

☐ the entire international application.

☒ claims Nos. 9,10.

because:

☐ the said international application, or the said claims Nos. relate to the following subject matter which does not require an international preliminary examination (*specify*):

☒ the description, claims or drawings (*indicate particular elements below*) or said claims Nos. 9,10 are so unclear that no meaningful opinion could be formed (*specify*):
see separate sheet

☐ the claims, or said claims Nos. are so inadequately supported by the description that no meaningful opinion could be formed.

☐ no international search report has been established for the said claims Nos. .

2. A meaningful international preliminary examination cannot be carried out due to the failure of the nucleotide and/or amino acid sequence listing to comply with the standard provided for in Annex C of the Administrative Instructions:

☐ the written form has not been furnished or does not comply with the standard.

☐ the computer readable form has not been furnished or does not comply with the standard.

V. Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)

Yes: Claims 1-8

**INTERNATIONAL PRELIMINARY
EXAMINATION REPORT**

International application No. PCT/GB00/03678

	No:	Claims	
Inventive step (IS)	Yes:	Claims	1-8
	No:	Claims	
Industrial applicability (IA)	Yes:	Claims	1-8
	No:	Claims	

2. Citations and explanations
see separate sheet

With respect to SECTION III:

Independent Claims 9 and 10 relate to subject-matter which is totally unclear. The subject-matter of these claims does *not* meet the requirements of Article 6 PCT (clarity) and in particular the stipulations of Rule 6.2 (a) PCT (references to other parts of the application) [see also PCT-Guidelines, PCT Gazette, Section IV, III-4.10]. These claims are therefore excluded from examination.

With respect to SECTION V:

1. The present international application PCT/GB00/03678 relates, according to the title, to packet network interfacing.
2. Claim 1 claims a method of establishing a *tunnel* from a first interface between a first network and a second network to a second interface between the second network and a third network.

Independent Claim 5 claims a method of sending packets from a first host on a first network via a second network to a second host on a third network

3. The nearest prior art is represented by the document
D1 = EP-A-0 840 482

This document D1 discloses, in accordance with features of Claim 1, a method of establishing a *tunnel* (cf. page 2, right-hand column, third to sixth paragraph; cf. fig. 19; cf. col. 21, starting on line 30 onwards) from a first interface (see fig. 19, "1") between a first network (see fig. 19, "106") and a second network (see fig. 19, "104") to a second interface (see fig. 19, "111") between the second network (see fig. 19, "104") and a third network (see fig. 19, "107"), the first (see fig. 19, "106") and third networks (see fig. 19, "107") operating in accordance with a first transmission protocol (IPv6 network) and having addresses in accordance with a first addressing convention (cf. col. 21, lines 44-48, e.g. "2::1"), and

the second network (see fig. 19, "104") operating in accordance with a **second** transmission protocol (IPv4 network) and having addresses in accordance with a **second** addressing convention (cf. col. 21, lines 44-48, e.g. "133.144.95.1"),
the *tunnel* being for the transport of messages from a first host (see fig. 19, "5") on the first network (106) to a second host (see fig. 19, "115") on the third network (107) (cf. page 2, right-hand column, third to sixth paragraph),

the method comprising the steps of:

- sending from the first host (5) an *address request message* (message Q; col. 22, line 7; fig. 20) in accordance with the first transmission protocol, referred to herein as a *first type address request message*, containing the **name** of the second host;
(cf. col. 22, lines 3-31; see fig. 20)
 - upon receipt of the address request message (message Q) at a name-to-address conversion system of the third network (see col. 22, lines 32-36), returning an *address response message* (message R; col. 22, line 37) in accordance with the **first** transmission protocol, referred to herein as a *first type address response message*, and containing the **address** of the second host in a response address field (see col. 22, lines 32-49).
4. The problems to be solved by the present international application are mainly explained on page 5 (line 20) to page 7 (line 14), on page 11 (lines 8 to 10) and on page 13 (lines 8 to 20). The main objectives of the present application are to render the transition from full IPv4 implementation (Internet Protocol version 4) to full IPv6 implementation possible and to find the true IPv6 address of a destination host.
- 5a. The problems are solved and the advantages are achieved with respect to Claim 1 by a method of establishing a *tunnel* from a first interface between a first network and a second network to a second interface between the second network and a third network, the first and third networks operating in accordance with a **first** transmission protocol and having addresses in accordance with a **first** addressing

convention, and

the second network operating in accordance with a **second** transmission protocol and having addresses in accordance with a **second** addressing convention,

the *tunnel* being for the transport of messages from a first host on the first network to a second host on the third network,

the method comprising the steps of:

- sending from the first host an *address request message* in accordance with the first transmission protocol, referred to herein as a *first type address request message*, containing the **name** of the second host;
- upon receipt of the address request message at a name-to-address conversion system of the third network,
returning an *address response message* in accordance with the first transmission protocol, referred to herein as a *first type address response message*, and containing the **address** of the second host in a response address field;
- upon receipt of that *first type address response message*
at the second interface,
converting it to an address response message in accordance with the **second** transmission protocol, referred to herein as a *second type address response message*, and
augmenting that converted *second type address response message* (42) by fields respectively containing
 - the address (50) of the second interface in accordance with the second addressing convention and
 - the address (52) of the second host in accordance with the first addressing convention; and
- upon receipt of that augmented that converted *second type address response message* **at the first interface**,
converting it to a *first type address response message*,
retrieving the contents of the augmenting fields,
storing at the first interface a mapping of the retrieved address of the second host and the retrieved address of the second interface for use in

encapsulating messages from the first host addressed to the second host,
and
replacing the content of the response address field of the resulting first type
address response message by the retrieved address of the second host.

- 5b. The problems are solved and the advantages are achieved
with respect to independent Claim 5
by a method of sending packets
from a first host on a first network via a second network
to a second host on a third network,
the first and third networks operating in accordance with a first transmission
protocol and having addresses in accordance with a first addressing
convention, and
the second network operating in accordance with a second transmission
protocol and having addresses in accordance with a second addressing
convention,
comprising the steps of:
- establishing a tunnel
from a first interface between the first network and the second network
to a second interface between the second network and the third network
in accordance with the method of claim 1;
 - upon receipt by the first host of the resulting first type address response
message from the first interface,
retrieving the content of the response address field;
 - generating one or more packets for transmission having a header
including source and destination address fields,
the source address field containing the address of the first host, and
the destination address field containing the retrieved content of the
response address field:
 - sending the or each generated packet to the first interface;

- for the or each said generated packet received by the **first interface**,
 - accessing the stored mappings in accordance with the destination address of the received generated packet,
 - retrieving the stored interface address of the mapping whose retrieved host address matches the destination address of the received generated packet,
 - generating an encapsulated packet having a payload formed by the received generated packet, and having a header including source and destination address fields, the source address field containing the address of the first interface, and the destination address field containing the retrieved interface network address,
 - sending the encapsulated packet to the second interface; and
 - for the or each encapsulated packet received by the **second interface**,
 - un-encapsulating the received encapsulated packet to recover the original generated packet forming its payload, and
 - sending that recovered packet to the second host.
6. Neither the above-mentioned document D1 nor the additionally cited documents of the international search report disclose all features of the respective independent claims. Therefore, the subject-matter defined in the above-mentioned independent claims 1 and 5 is novel (Article 33 (2) PCT).
Moreover, none of the documents cited, neither considered alone nor in combination with another document, is directed to the problem to be solved or leads a skilled person in an obvious manner to the combination of technical features as claimed in the independent claims. Thus, the subject-matter defined in the above-mentioned independent claims involves an inventive step (Article 33 (3) PCT).
7. Dependent Claims 2 to 4 and 6 to 8 relate to further details of the methods and are therefore equally novel and inventive (Art. 33 (2) and (3) PCT).
8. Industrial applicability (Article 33 (4) PCT) of the subject-matter claimed is beyond doubt.

With respect to clarity of independent Claim 5:

The subject-matter of Claim 5 includes the specification

"- establishing a tunnel
from a first interface between the first network and the second network
to a second interface between the second network and the third network
in accordance with the method of claim 1; "

The reference to the method of Claim 1 could render the wording of Claim 5 unclear. This is because it could be ambiguous which features of Claim 1 are exactly meant by "establishing a tunnel ... in accordance with the method of claim 1".

It should be noted that clarity (Art. 6 PCT) is of utmost importance for the purposes of formulating an opinion on the questions whether the claimed invention appears to be novel, to involve an inventive step and to be industrially applicable in view of their function in defining the matter for which protection is sought (cf. PCT Gazette, Section IV, III-4.1). Moreover, the subject-matter claimed should be clear from the wording of the claim alone (cf. PCT Gazette, Section IV, III-4.2).

Therefore, the specification "in accordance with the method of claim 1" should be replaced by the method steps meant and defined in Claim 1.

Formal aspects which should be considered during the following patent procedures:

- 1). To meet the requirements of Rule 6.3 (b) PCT, the independent claims should be properly cast in the two-part form, with those features which in combination are part of the prior art (D1) being placed in the preamble.
- 2). Reference signs in parentheses should be inserted in the claims to increase their intelligibility, Rule 6.2 (b) PCT. This applies to both the preamble and the characterising portion.
- 3). To meet the requirements of Rule 5.1 (a)(ii) PCT, the relevant document D1 of the international search report should be identified in the description and the relevant background art disclosed therein should be briefly discussed.

- 4). On page 1, lines 11 to 13, the text reads: *"Herein, the terms packet and message are used interchangeably and have the same meaning, and the term Internet domain is used as a specific example of a network."*
- Moreover, on page 9, lines 28-30, it is stated that *"For convenience, the terms 'field' and 'record' are used synonymously and interchangeably in this description, although in the art a field is generally taken to be a component of a record."*

It should be noted that the stipulations of Rule 10.2 PCT require that the terminology and the signs shall be consistent throughout the international application. Hence, the text quoted above is not in line with the requirements of Rule 10.2 PCT (as well as to clarity). The reader is confused when studying the application if the terms *"packet"* and *"message"* or *"field"* and *"record"* have the same meaning and if these terms are used interchangeably.

Furthermore, the problems to be solved specifically concern the Internet. Hence, a generalization of the problems, and the solution to other (non-Internet) networks, does not seem to be possible.

2. On page 10, lines 1 to 4, it is stated that *"The controller 64B is **connected via control line 65A** to control inputs of ... and **the 6to4 encapsulator 90B** ... (emphasis added)"*. However, in figure 2, there is no connection between the control line 65 and the 6to4 encapsulator 90.
- Moreover, figure 2 depicts a general structure of a border router or an interface 16. This structure is respectively used in the interfaces 16A and 16B shown in figure 1. Hence, the extension "A" together with a reference sign shown in figure 2 refers to a component of the interface 16A and the extension "B" together with a reference sign shown in figure 2 refers to a component of the interface 16B.
- The above-cited passage on page 10, however, refers to a control line 65A (meaning that the control line is part of interface 16A) which is connected to parts 64B, 32B, 86B and 90B being parts of the interface 16B. It seems impossible that the control line 65A is contained in the interface 16B (cf. also e.g. page 10, lines 10 and 32).

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference A25800 WO	FOR FURTHER ACTION <small>see Notification of Transmittal of International Search Report (Form PCT/ISA/220) as well as, where applicable, item 5 below.</small>	
International application No. PCT/GB 00/ 03678	International filing date (<i>day/month/year</i>) 25/09/2000	(Earliest) Priority Date (<i>day/month/year</i>) 24/09/1999
Applicant BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY		

This International Search Report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This International Search Report consists of a total of 3 sheets.



It is also accompanied by a copy of each prior art document cited in this report.

1. Basis of the report

- a. With regard to the **language**, the international search was carried out on the basis of the international application in the language in which it was filed, unless otherwise indicated under this item.



the international search was carried out on the basis of a translation of the international application furnished to this Authority (Rule 23.1(b)).

- b. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, the international search was carried out on the basis of the sequence listing :



contained in the international application in written form.



filed together with the international application in computer readable form.



furnished subsequently to this Authority in written form.



furnished subsequently to this Authority in computer readable form.



the statement that the subsequently furnished written sequence listing does not go beyond the disclosure in the international application as filed has been furnished.



the statement that the information recorded in computer readable form is identical to the written sequence listing has been furnished

2. ☐ **Certain claims were found unsearchable** (See Box I).

3. ☐ **Unity of invention is lacking** (see Box II).

4. With regard to the **title**,



the text is approved as submitted by the applicant.



the text has been established by this Authority to read as follows:

5. With regard to the **abstract**,



the text is approved as submitted by the applicant.



the text has been established, according to Rule 38.2(b), by this Authority as it appears in Box III. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.

6. The figure of the **drawings** to be published with the abstract is Figure No.



as suggested by the applicant.



because the applicant failed to suggest a figure.



because this figure better characterizes the invention.

2



None of the figures.

INTERNATIONAL SEARCH REPORT

International Application No

/GB 00/03678

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

PAJ, EPO-Internal, IBM-TDB, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 802 285 A (HIRVINIEMI SEPPÖ) 1 September 1998 (1998-09-01) abstract page 2, right-hand column, line 13 -page 3, left-hand column, line 28 ----	1-10
A	LABORDE D: "UNDERSTANDING AND IMPLEMENTING EFFECTIVE VPNS" COMPUTER TECHNOLOGY REVIEW, US, WESTWORLD PRODUCTION CO. LOS ANGELES, vol. 18, no. 2, 1 February 1998 (1998-02-01), page 12, 14, 16 XP000733729 ISSN: 0278-9647 the whole document ----- -/--	1-10



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

20 November 2000

Date of mailing of the international search report

28/11/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Perez Perez, J

INTERNATIONAL SEARCH REPORT

International Application No

/GB 00/03678

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 840 482 A (HITACHI LTD) 6 May 1998 (1998-05-06) claims	1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

/GB 00/03678

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
US 5802285	A	01-09-1998	FI	90710 B	30-11-1993
			GB	2267418 A, B	01-12-1993
EP 0840482	A	06-05-1998	JP	10136052 A	22-05-1998
			JP	11055319 A	26-02-1999
			US	6118784 A	12-09-2000